

AFRL-VA-WP-TP-2002-303

**DYNAMIC INVERSION-BASED
ADAPTIVE/RECONFIGURABLE
CONTROL OF THE X-33 ON ASCENT**

**David B. Doman
Anh Tuan D. Ngo**



APRIL 2002

Approved for public release; distribution is unlimited.

This material is declared a work of the U.S. Government and is not subject to copyright protection in the United States.

**AIR VEHICLES DIRECTORATE
AIR FORCE RESEARCH LABORATORY
AIR FORCE MATERIEL COMMAND
WRIGHT-PATTERSON AIR FORCE BASE, OH 45433-7542**

20020912 081

REPORT DOCUMENTATION PAGE

Form Approved
OMB No. 0704-0188

The public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Department of Defense, Washington Headquarters Services, Directorate for Information Operations and Reports (0704-0188), 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to any penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number. PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ADDRESS.

1. REPORT DATE (DD-MM-YY) April 2002			2. REPORT TYPE Journal Article		3. DATES COVERED (From - To) 06/01/1999 - 10/01/2001	
4. TITLE AND SUBTITLE DYNAMIC INVERSION-BASED ADAPTIVE/RECONFIGURABLE CONTROL OF THE X-33 ON ASCENT					5a. CONTRACT NUMBER In-house	
					5b. GRANT NUMBER	
					5c. PROGRAM ELEMENT NUMBER 61102F	
6. AUTHOR(S) David B. Doman Anh Tuan D. Ngo					5d. PROJECT NUMBER 2304	
					5e. TASK NUMBER N2	
					5f. WORK UNIT NUMBER 01	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Control Theory Optimization Branch (AFRL/VACA) Control Sciences Division Air Vehicles Directorate Air Force Research Laboratory, Air Force Materiel Command Wright-Patterson Air Force Base, OH 45433-7542					8. PERFORMING ORGANIZATION REPORT NUMBER AFRL-VA-WP-TP-2002-303	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES) Air Vehicles Directorate Air Force Research Laboratory Air Force Materiel Command Wright-Patterson AFB, OH 45433-7542					10. SPONSORING/MONITORING AGENCY ACRONYM(S) AFRL/VACA	
					11. SPONSORING/MONITORING AGENCY REPORT NUMBER(S) AFRL-VA-WP-TP-2002-303	
12. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release; distribution is unlimited.						
13. SUPPLEMENTARY NOTES This material is declared a work of the U.S. Government and is not subject to copyright protection in the United States. Article published in <i>Journal of Guidance, Control, and Dynamics</i> , March-April 2002 (Vol. 25, No. 2).						
14. ABSTRACT A quaternion-based attitude control system is developed for the X-33 in the ascent flight phase. A nonlinear control law commands body-axis rotation rates that align the angular velocity vector with an Euler axis defining the axis of rotation that will rotate the body-axis system into a desired-axis system. The magnitudes of the commanded body rates are determined by the magnitude of the rotation error. The commanded body rates form the input to a dynamic inversion-based adaptive/reconfigurable control law. The indirect adaptive control portion of the control law uses online system identification to estimate the current control effectiveness matrix to update a control allocation module. The control allocation nominally operates in a minimum deflection mode; however, if a fault is detected, it can operate in a null-space injection mode that excites and decorrelates the effectors without degrading the vehicle response to enable online system identification. The overall system is designed to provide fault and damage tolerance for the X-33 on ascent.						
15. SUBJECT TERMS						
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT: SAR	18. NUMBER OF PAGES 18	19a. NAME OF RESPONSIBLE PERSON (Monitor) Dr. David B. Doman	
a. REPORT Unclassified	b. ABSTRACT Unclassified	c. THIS PAGE Unclassified			19b. TELEPHONE NUMBER (Include Area Code) (937) 255-8451	

Dynamic Inversion-Based Adaptive/ Reconfigurable Control of the X-33 on Ascent

D. B. Doman and A. D. Ngo

Reprinted from

Journal of Guidance, Control, and Dynamics

Volume 25, Number 2, Pages 275-284



A publication of the
American Institute of Aeronautics and Astronautics, Inc.
1801 Alexander Bell Drive, Suite 500
Reston, VA 20191-4344

Dynamic Inversion-Based Adaptive/Reconfigurable Control of the X-33 on Ascent

David B. Doman* and Anhtuan D. Ngo†

U.S. Air Force Research Laboratory, Wright–Patterson Air Force Base, Ohio 45433-7531

A quaternion-based attitude control system is developed for the X-33 in the ascent flight phase. A nonlinear control law commands body-axis rotation rates that align the angular velocity vector with an Euler axis defining the axis of rotation that will rotate the body-axis system into a desired-axis system. The magnitudes of the commanded body rates are determined by the magnitude of the rotation error. The commanded body rates form the input to a dynamic inversion-based adaptive/reconfigurable control law. The indirect adaptive control portion of the control law uses online system identification to estimate the current control effectiveness matrix to update a control allocation module. The control allocation nominally operates in a minimum deflection mode; however, if a fault is detected, it can operate in a null-space injection mode that excites and decorrelates the effectors without degrading the vehicle response to enable online system identification. The overall system is designed to provide fault and damage tolerance for the X-33 on ascent. The baseline control law is based on a full envelope design philosophy and eliminates trajectory-dependent gain scheduling that is typically found on this type of vehicle. Results are shown to demonstrate the feasibility of the approach.

Nomenclature

$A1, A2,$	=	aerospike engine quadrants
$B1, B2$		
a	=	command limiting gain growth/decay constant
B	=	control effectiveness matrix
C	=	direction cosine
\dot{d}_{des}	=	desired control variable rate
e	=	eigenvector
F	=	force vector
G	=	moment vector
H	=	regressor matrix containing measured control deflections over finite horizon
I	=	inertia tensor
J	=	cost function
K_B	=	bandwidth of desired dynamics
K_p	=	principal attitude error gain
L, M, N	=	roll, pitch, and yaw moments
P, Q, R	=	covariance matrices
P	=	parameter vector
p, q, r	=	roll, pitch, and yaw velocities
q	=	scalar quaternion component
\bar{q}	=	vector part of quaternion
\bar{q}	=	quaternion
R	=	rotation matrix
s	=	Laplace transform operator
v	=	vector of uniformly distributed random variables
\bar{W}	=	diagonal scaling matrix for command distribution
W_δ	=	weighting vector on control deflections
w, W	=	measurement noise
Z	=	vector of measured accelerations over finite horizon
α, β	=	angle of attack, sideslip angle
ΔT	=	flight control system update rate
δ	=	control effector deflection vector
$\delta, \bar{\delta}$	=	most/least restrictive control range limits
δ_p	=	control preference vector
δ_r	=	vector of effector rate limits

δ_s	=	slack variable vector
Θ	=	parameter vector to be estimated
ϕ	=	principal attitude error
ω	=	vehicle angular velocity vector
$\ \bullet\ _1$	=	1-norm of a vector
$\hat{\bullet}$	=	estimate

I. Introduction

As a subscale technology demonstrator for a quick-turnaround, reusable launch vehicle, the X-33 (Fig. 1) is an autonomous, hypersonic vehicle with two linear aerospike engines and eight aerodynamic control surfaces: inner/outer elevons, rudders, and flaps. The vehicle takes off vertically like a rocket and lands horizontally like an airplane. The unpiloted vehicle has a flight envelope that spans up to 30 miles in altitude and Mach 9 in speeds.

Adaptive/reconfigurable flight control technologies have been maturing over the past decade. A number of different approaches have been developed, and some have been tested on fighter-type aircraft. An indirect adaptive control approach¹ was demonstrated on the VISTA-F-16 in 1995, where a simulated failure of a left horizontal tail was induced on approach and the vehicle landed without incident. More recently, the X-36 tailless remotely piloted vehicle successfully demonstrated a direct adaptive control system where a neural network was used to regulate adaptively the inversion error of a baseline dynamic inversion control law.² The U.S. Air Force Reconfigurable Systems for Tailless Fighter Aircraft program developed a number of direct and indirect adaptive control algorithms for tailless fighter aircraft.^{3–6} In general, indirect adaptive control systems require online identification of the model parameters that are used for the online computation of a control law. Direct adaptive control schemes do not require explicit identification of model parameters; instead, they generate signals that attempt to cancel the effects of modeling error on control laws that were designed using nominal or unfaulted models of the controlled element.

Transitioning adaptive control technology to reusable launch vehicles (RLVs) such as the X-33 has the potential of providing benefits demonstrated on fighter aircraft to this new class of vehicle. Benefits include improved reliability through fault and damage tolerance to effector failures and faster turnaround times through reduced flight control redesign times. The current X-33 control system is essentially a trajectory dependent gain-scheduled proportional-integral-derivative architecture.⁷ The system is reconfigurable in that it is designed to accommodate a predefined set of failures such as engine-out, again using gain scheduling. Adaptive/reconfigurable control is

Received 18 January 2001; revision received 27 August 2001; accepted for publication 4 September 2001. This material is declared a work of the U.S. Government and is not subject to copyright protection in the United States. Copies of this paper may be made for personal or internal use, on condition that the copier pay the \$10.00 per-copy fee to the Copyright Clearance Center, Inc., 222 Rosewood Drive, Danvers, MA 01923; include the code 0731-5090/02 \$10.00 in correspondence with the CCC.

*Aerospace Engineer, Air Vehicles Directorate. Senior Member AIAA.

†Electronics Engineer, Air Vehicles Directorate. Member AIAA.




Fig. 1 X-33 RLV.

different in that the system adapts to failures (possibly unforeseen) online in real time.

This paper will discuss an indirect adaptive control system that is built around a baseline dynamic inversion control law. Dynamic inversion control laws require the use of a control mixer or control effector allocation algorithm when the number of control effectors exceeds the number of controlled variables. This is because a small number of desired moment or acceleration commands are calculated and a large number of control effectors may be used to achieve the desired command. It is quite common that the desired commands can be achieved in many different ways, and so control allocation approaches are used to provide consistent and unique solutions to such problems. The control allocation relies on accurate knowledge of the control derivatives. Under failure or damage conditions, the control derivatives can be altered dramatically. When control derivatives are identified and updated information is supplied to the control allocation block, the performance of the entire system can be improved. A fault-detection scheme is used to trigger an online system identification algorithm that is used to estimate control effectiveness when faults are detected. The control system described in this paper is implemented in C and interfaced with the MAVERIC simulation of the X-33.

II. Baseline Attitude Control System

During the ascent flight phase, the X-33 control system must track attitude commands generated by the guidance system. The use of a dynamic inversion control law to control the Euler angles of this vehicle on ascent is precluded because of the requirement to compute the Euler angle rates. The pitch attitude of the vehicle changes by more than 90 deg throughout this flight phase, which brings the well-known singularity problems inherent in the computation of Euler angle rates into play. To circumvent this problem, a quaternion-based control law is developed and coupled with a dynamic inversion-based body-axis rate command system. The system works by determining the axis of rotation that will take the vehicle's body-axis system directly to a desired-axis system in a single rotation. This axis is called an Euler axis, and all vectors that lie on it have an interesting property in that they have the same components in the body-axis system as they do in the desired-axis system. This characteristic arises from all rotation matrices R having one eigenvalue equal to 1 and the Euler axis being defined by the eigenvector associated with this eigenvalue; thus, $Re_1 = e_1$. Therefore, it is easy to see that the Euler axis has the same components in the transformed system as it does in the original coordinate frame. A set of body-axis rate commands is then generated that are aligned with a body-axis to desired Euler axis. The magnitude of the commanded rates is a function of the magnitude of the scalar angular misalignment ϕ between the body- and desired-axis systems.

In general we define a quaternion as having scalar and vector parts

$$\tilde{q} = q_0 + \mathbf{q} = q_0 + [q_1 \ q_2 \ q_3] \quad (1)$$

with a constraint on the unit norm

$$q_0^2 + q_1^2 + q_2^2 + q_3^2 = 1 \quad (2)$$

A body-to-desired quaternion \tilde{q}_{b2d} holds information about the direction of the Euler axis in its vector part \mathbf{q}_{b2d} and information about the angular misalignment of the two axis systems in its scalar part q_{0b2d} . The scalar and vector parts of the body-to-desired quaternion are given by

$$q_{0b2d} = \cos(\phi/2), \quad \mathbf{q}_{b2d} = \begin{bmatrix} C_1 \sin(\phi/2) \\ C_2 \sin(\phi/2) \\ C_3 \sin(\phi/2) \end{bmatrix} \quad (3)$$

where C_1 , C_2 , and C_3 are the direction cosines of the Euler axis relative to the body-axis reference frame.

The guidance system generates attitude commands referenced with respect to a plumbline-axis system. The plumbline-axis system is a vehicle carried frame that is attached to the vehicle center of gravity, whose z axis points toward the center of the Earth and whose x axis is oriented along the launch azimuth with the y axis completing the right-handed coordinate system. Two quaternions are used to compute the body quaternion to desired quaternion, the plumbline quaternion to desired quaternion \tilde{q}_{p2d} and the plumbline quaternion to body quaternion \tilde{q}_{p2b} . The body quaternion to desired quaternion as defined can be calculated using quaternion multiplication and inversion rules⁸:

$$\tilde{q}_{b2d} = \tilde{q}_{p2d} \tilde{q}_{p2b}^{-1} \quad (4)$$

The vector part of \tilde{q}_{b2d} yields the direction of the Euler axis in body-axis coordinates. Quaternion multiplication and inversion is a more computationally tractable approach that can be used to determine the Euler axis and magnitude of the alignment error when compared to direction cosine matrices because the online computation of eigenvalues and eigenvectors and the associated sign ambiguities are side stepped.

Figure 2 shows a block diagram of the baseline ascent attitude control system that was developed in this work. The motivation for the design is that, if at any instant, the body-axis angular velocity vector is aligned with the body axis to desired Euler axis, the body-axis system will move toward the desired-axis system thereby reducing the magnitude of the attitude error. When the attitude error ϕ is zero, $q_{0b2d} = 1$ and $\mathbf{q}_{b2d} = [0 \ 0 \ 0]$ as can be deduced from Eq. (3).

Note from Fig. 2 that a feasibility test result from the control allocation block can modify the principle attitude-error gain. As will be discussed later, a two-branch control allocation scheme is used that either minimizes the difference between the desired and attainable body-axis accelerations or drives the effectors to preferred positions when sufficient control power exists. The former case is referred to as a control deficiency condition, and it occurs when the desired accelerations are unattainable due to effector constraints such as rate or position limits. Typically one axis will saturate first, which can produce angular velocity vectors that are not aligned with the Euler axis. To prevent this condition, the principle attitude error gain K_p is reduced when any one of the commanded accelerations becomes infeasible to reduce collectively the angular velocity commands. The gain K_p is bounded by $0 \leq K_p \leq K_{p, \text{nom}}$, where $K_{p, \text{nom}}$ is the nominal value of K_p . The gain K_p decreases exponentially toward zero when the command is infeasible and increases exponentially toward $K_{p, \text{nom}}$ when the command is feasible. The rate of decay or growth is governed by the choice of the constant a , which was set to unity in this application.

III. Dynamic Inversion

The quaternion-based outer-loop control system generates body-axis angular velocity commands p_c , q_c , and r_c that are aligned with the error Euler axis. The inner-loop dynamic inversion control law is designed so that the X-33 tracks these body-rate commands. The X-33 rotational dynamics can be written as

$$\dot{\omega} = f(\omega, \mathbf{P}) + \mathbf{g}(\mathbf{P}, \delta) \quad (5)$$

where $\omega = [p \ q \ r]$ and \mathbf{P} denotes measurable or estimable quantities that influence the body-rate states. The parameter vector \mathbf{P} includes variables such as Mach number, angle of attack, sideslip angle, and vehicle mass properties such as moments of inertia. Equation (5) expresses the body-axis rotational accelerations as a sum that includes control-dependent accelerations $\mathbf{g}(\mathbf{P}, \delta)$ and accelerations that are due only to the base engine and aerodynamics.

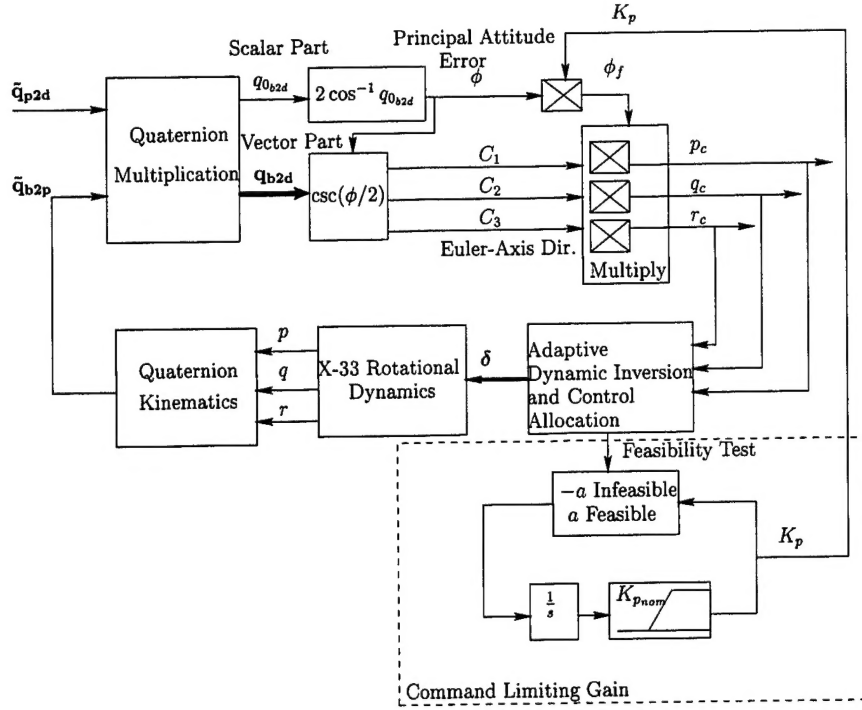


Fig. 2 Conceptual block diagram of quaternion-based attitude control system for ascent.

It is assumed that the mass properties of the X-33 change slowly when compared to the body-axis rates so that $\dot{I} \approx 0$ and

$$\dot{\omega} = I^{-1}(G_B - \omega \times I\omega) \quad (6)$$

where

$$G_B = G_{BAE}(\omega, P) + G_\delta(P, \delta) = \begin{bmatrix} L \\ M \\ N \end{bmatrix}_{BAE} + \begin{bmatrix} L \\ M \\ N \end{bmatrix}_\delta \quad (7)$$

where $G_{BAE}(\omega, P)$ is the moment generated by the base aerodynamics and engine (BAE) system and G_δ is the sum of the moments produced by the control effectors. Thus,

$$f(\omega, P) = I^{-1}[G_{BAE}(\omega, P) - \omega \times I\omega] \quad (8)$$

$$g(P, \delta) = I^{-1}G_\delta(P, \delta) \quad (9)$$

Dynamic inversion requires that the control-dependent portion of the model be affine in the controls. Therefore, we develop a linear approximation of the control-dependent part such that

$$G_\delta(P, \delta) \approx G_\delta(P)\delta \quad (10)$$

The X-33 aerodynamic database provides force and moment coefficient data that are taken at a moment reference point (MRP), which is located at the center of gravity of the empty vehicle, that is, no fuel/oxidizer. Control derivative information was extracted from the tables in the database for Mach numbers, angles of attack, and sideslip angles that were to be encountered on the ascent trajectory. Polynomial fits to the discrete control derivative data were produced to provide continuous estimates of $G_\delta(P)$. The control derivatives are continuously corrected for the moving center of gravity as the vehicle ascends using the following relation:

$$G_\delta(P) = G_{\delta_{MRP}} + (r_{MRP} - r_{cg}) \times F_{\delta_{MRP}} \quad (11)$$

The X-33 power pack is an XRS 2200 linear aerospike rocket engine. This engine is divided into four sections: port upper A1, port lower A2, starboard upper B1, and starboard lower B2. These quadrants can be differentially throttled to generate moments that

can be used for attitude control. This differential throttling is critical at launch because the aerodynamic surfaces are ineffective at low dynamic pressure. The engine control derivatives are estimated by using a global slope approximation that simply divides the estimated torque produced by a quadrant by the actual chamber pressure in that quadrant. For example, the torque gradient of the upper port quadrant A1 would be estimated by

$$\frac{\partial G_{A1}}{\partial P_{cA1}} \approx \frac{G_{A1}}{P_{cA1}} \quad (12)$$

The engine control derivatives are also corrected for the moving c.g. The engine torques G_{A1} , G_{A2} , G_{B1} , and G_{B2} can be estimated by feeding chamber pressure commands through a transfer function model of the engine with limits to estimate the actual chamber pressures. The engine forces and moments can then be calculated using a table lookup model parameterized by chamber pressure, mixture ratio, and pressure ratio. This process was used to determine a single average estimate of the engine control derivatives. This simple model was sufficient to provide good tracking performance over a very wide range of flight conditions.

The model used for the design of the dynamic inversion control law becomes

$$\dot{\omega} = f(\omega, P) + G_\delta(P)\delta \quad (13)$$

and our objective is to find a control law that provides direct control over $\dot{\omega}$ so that $\dot{\omega} = \dot{\omega}_{des}$, that is,

$$\dot{\omega}_{des} = f(\omega, P) + G_\delta(P)\delta \quad (14)$$

Therefore, the inverse control must satisfy

$$\dot{\omega}_{des} - f(\omega, P) = G_\delta(P)\delta \quad (15)$$

Because there are more control effectors than controlled variables, a control allocation algorithm must be used to obtain a unique solution. Control allocation will be discussed in detail in the next section. There are 12 control effectors that may be used on ascent: inboard and outboard elevons, left and right rudders, body flaps, and chamber pressures of the four quadrants of the aerospike engine. Equation (15) states that the control effectors are to be used to correct for the difference between the desired accelerations and




Fig. 3 Block diagram of inner-loop dynamic inversion control law.




Fig. 4 Implicit model-following implementation using prefilter and dynamic inversion.

the accelerations due only to the base engine and aerodynamic moments.

When the loop is closed around the aircraft, as shown in Fig. 3, and sufficient control power exists such that Eq. (15) is satisfied, the transfer function matrix should approach a bank of decoupled integrators:

$$\begin{bmatrix} p \\ q \\ r \end{bmatrix} \approx \begin{bmatrix} 1/s & 0 & 0 \\ 0 & 1/s & 0 \\ 0 & 0 & 1/s \end{bmatrix} \begin{bmatrix} \dot{p}_{des} \\ \dot{q}_{des} \\ \dot{r}_{des} \end{bmatrix} \quad (16)$$

The higher the fidelity of the model used in the dynamic inversion control law, the more the closed-loop system will behave like a decoupled bank of integrators.

A command-shaping prefilter is used to convert commands from the outer-loop attitude controllers p_c , q_c , and r_c into acceleration commands to the dynamic inversion control law \dot{p}_{des} , \dot{q}_{des} , and \dot{r}_{des} . The prefilter structure is designed to provide a first-order low-pass filter response for each command variable (CV) (p , q , r) to desired command variable signal (CV^{CMD}) (p_c , q_c , r_c). The combination of prefilter and dynamic inversion produces an implicit model-following framework. Figure 4 shows a block diagram of the prefilter that is used for each axis.

The gain K_B can be interpreted as the crossover frequency of the loop transfer function. Achieving the desired closed-loop dynamics is critically dependent on the dynamic inversion/control allocation algorithm successfully producing a decoupled bank of integrators. If dynamic inversion is perfect the closed-loop transfer function for each command variable becomes

$$\frac{CV}{CV^{CMD}} = \frac{\frac{1}{2} K_B}{s + \frac{1}{2} K_B} \quad (17)$$

The gain K_B must be selected to provide sufficiently high bandwidth tracking without over-driving the control effectors. If the dynamic inversion and control allocation are perfect, Eq. (17) could be achieved using a simple gain on the tracking error. Unfortunately, there always exists some modeling error, and the actuators cannot respond instantaneously to the rate- and position-limited command signals generated by the control allocation algorithm. The more complicated prefilter structure provides an integral action on the error to compensate for modeling errors and to improve steady-state tracking error while the feedforward structure balances the lag introduced by the integral compensation to maintain phase margin near crossover.

In summary, the fundamental objective of the dynamic inversion control law is to provide good body angular rate tracking. The dynamic inversion and model-following architectures will be augmented with direct and indirect adaptive control algorithms to mitigate the impact of uncertainties and to compensate for damage and failures.

IV. Control Allocation

There are 3 controlled variables on the ascent trajectory and 12 control effectors; therefore, a control allocation scheme must be used to ensure that Eq. (15) is satisfied. The control allocation scheme used in this case draws heavily on the work of Buffington⁴ and Buffington et al.⁹ The control allocation problems are formulated as linear programs (LPs). The LP approach minimizes a linear performance index subject to linear constraints. Linear inequality constraints are used to ensure that effector rate and position limits are not violated. More complex engine constraints are also accommodated to ensure that feasible thrust vectoring commands are generated. The control allocation problem is broken down into a control deficiency branch and a control sufficiency branch.

Numerous codes exist to solve such problems. The widely used LP.Solve package written in C programming language was used in this case to be compatible with the X-33 MAVERIC simulation and was found to perform well in many cases; however, a number of Monte Carlo dispersion runs performed by a NASA Marshall Space Flight Center evaluation team uncovered several cases where the LP.Solve algorithm failed to converge. Up to 20,000 linear programming problems are posed and solved over the course of a typical ascent trajectory. Occasionally (approximately 1 in 100,000) one of these problems is degenerate and fails to converge. Interestingly, these degenerate problems can be solved by some linear programming packages and not others. At the present time, there are obviously serious flight certification and implementation issues associated with optimization-based online control allocation, but the benefits are great, and these issues will undoubtedly be resolved in time. Some of the solvers that have been capable of solving these degenerate problems remain to be integrated into the MAVERIC simulation. New fast and reliable LP-based online control allocation methods are being developed¹⁰ and may also eliminate these deficiencies.

Control Deficiency Branch

The control deficiency branch is used to test feasibility of satisfying Eq. (15). For convenience we will refer to the left-hand side of Eq. (15) as d_{des} :

$$d_{des} \triangleq \dot{\omega}_{des} - f(\omega, P) = G_\delta(P) \delta \triangleq B\delta \quad (18)$$

If it is not feasible to obtain $d_{des} = B\delta$ due to control effector constraints, then the difference between the desired and actual effector-induced body-axis accelerations is minimized. Thus, the objective can be summarized in terms of minimizing a 1-norm performance index subject to constraints:

$$\min_{\delta} J_D = \|B\delta - d_{des}\|_1, \quad \text{subject to } \underline{\delta} \leq \delta \leq \bar{\delta} \quad (19)$$

where $\underline{\delta}$ and $\bar{\delta}$ are the most restrictive lower bounds and upper bounds on the control effector deflection. Here,

$$\bar{\delta} = \min(\delta_u, \Delta T \dot{\delta}_r + \delta), \quad \underline{\delta} = \max(\delta_l, -\Delta T \dot{\delta}_r + \delta) \quad (20)$$

where δ_u is the upper position limit vector, δ_l is the lower position limit vector, and ΔT is the inner-loop flight control system update rate. The optimization problem posed in Eq. (19) may be transformed into the following linear programming problem:

$$\begin{aligned} \min_{\delta} J_D &= [0 \quad \dots \quad 0 \quad 1 \quad \dots \quad 1] \begin{bmatrix} \delta \\ \delta_s \end{bmatrix} \\ \text{subject to } &\begin{bmatrix} -\delta_s \\ \delta \\ -\delta \\ B\delta - \delta_s \\ -B\delta - \delta_s \end{bmatrix} \leq \begin{bmatrix} 0 \\ \bar{\delta} \\ -\underline{\delta} \\ d_{des} \\ -d_{des} \end{bmatrix} \end{aligned} \quad (21)$$

where δ_s is of the same dimension as the set of controlled variables. If $J_D = 0$, then the commanded controlled variable rates are achievable, and there may be excess control power available that can

be used to optimize subobjectives. If $J \neq 0$, the commanded controlled variable rates are not achievable, and the control allocation algorithm provides a vector of effector commands that minimizes the deficiency.

Control Sufficiency Branch

If there is sufficient control power available such that $J_D = 0$, then there may be excess control power available to optimize a subobjective. The subobjective could involve driving the control effectors to a preferred position δ_p . A performance index reflecting this objective is given by

$$\min_{\delta} J_S = \|W_{\delta}(\delta - \delta_p)\|_1$$

subject to $B\delta = d_{des}, \quad \underline{\delta} \leq \delta \leq \bar{\delta}$ (22)

where W_{δ} is a vector that allows one to weight one preference over another. This optimization problem can be cast into the LP framework as follows:

$$\min_{\delta} J_S = W_{\delta}^T \delta_s$$

subject to $\begin{bmatrix} -\delta_s \\ \delta \\ -\delta \\ \delta - \delta_s \\ -\delta - \delta_s \end{bmatrix} \leq \begin{bmatrix} 0 \\ \bar{\delta} \\ -\underline{\delta} \\ \delta_p \\ -\delta_p \end{bmatrix}, \quad B\delta = d_{des}$ (23)

where δ , δ_s , δ_p , and W_{δ} are of the same dimension as the number of control effectors. The preference vector δ_p is used in this case to decorrelate the control effectors to enable online system identification of the control effectiveness matrix B .

Null-Space Injection

The indirect adaptive portion of the control law requires online identification of the control effectiveness matrix B . This enables the control law to compensate for failures, damage, or modeling errors throughout the flight. To identify elements of the control effectiveness matrix, each control effector must be active at all times. Furthermore, each effector must be moving independently so that there is no correlation between the movement of one control effector and another. Decorrelated control deflections are necessary to obtain a well-conditioned regressor matrix for system identification. One way of doing this is to provide dithered effector commands that consist of an additive random signal that is superimposed on the nominal effector command. Unfortunately, this simple approach results in degraded vehicle response because in general $B(\delta + \delta_{dither}) \neq d_{des}$. The solution is to provide a dithering signal that lies in the null space of the B , that is, $B\delta_{dither} = 0$, so that $B(\delta + \delta_{dither}) = d_{des}$. This can be accomplished indirectly by randomly perturbing the control effector preference vector according to

$$\delta_p \triangleq W^{-1}B^T(BW^{-1}B^T)^{-1}d_{des}$$
 (24)

where

$$W = \bar{W}W_r, \quad W_r = \text{diag}(10^{v_1}, 10^{v_2}, \dots, 10^{v_m})$$
 (25)

and v is a vector of uniformly distributed random variables between -1 and 1 . The matrix \bar{W} is a nominal diagonal weighting matrix used for scaling purposes to distribute commands equally. Note that δ_p is actually the solution to a weighted least-squares problem:

$$\min_{\delta} J = \delta^T W \delta, \quad \text{subject to } B\delta = d_{des}$$
 (26)

Thus, the preference vector will be driven toward a randomly weighted least-squares solution to the control allocation problem that does not account for rate and position constraints. Now the preference vector δ_p is randomly changing, and the sufficiency branch of the LP-based control allocation ensures that $B\delta = d_{des}$ and that the control effector constraints are not violated. This approach ensures




Fig. 5 Attitude tracking using dynamic inversion/null-space injection control allocation.




Fig. 6 Attitude tracking using dynamic inversion/minimum deflection control allocation.

that the control effectors are decorrelated and active without degrading the vehicle response. This approach also avoids the explicit calculation of the null space of B .

Figure 5 shows the response of the body flaps and the tracking performance of the X-33 operating in the null-space injection mode. One can see that the flaps are quite active and decorrelated while the tracking performance is well behaved. Figure 6 shows the flap response and tracking performance of the X-33 operating with minimum-deflection control allocation. The differences between the null-space injection mode and the minimum deflection mode can easily be seen. In the minimum deflection mode, some of the effectors are completely inactive, for example, rudders, while others are highly correlated, for example, flaps. The vehicle tracking performance is good as well. In general it is more advantageous to operate in the minimum deflection mode due to the lower energy requirements; therefore, the null-space injection mode is only turned on when a fault is detected and system identification is required. The fault detection scheme will be discussed later.




Fig. 7 Attitude tracking using dynamic inversion/null-space injection control allocation.




Fig. 8 Body-rate tracking using dynamic inversion/null-space injection control allocation.

V. Baseline Dynamic Inversion Results

The baseline attitude control system shown in Figs. 2–4 was implemented in the MAVERIC simulation of the X-33. The desired roll, pitch, and yaw bandwidths were set to $K_b = 5$, and the principal attitude error proportional gain was set to $K_p = 1$. Null-space injection control allocation was implemented using LP techniques and was used to excite and decorrelate the control effectors without degrading the vehicle response. The commanded and actual Euler angles for the ascent trajectory are shown in Fig. 7. One can see that the X-33 tracks the commanded attitude quite well. The stepped nature of the commanded attitude is a result of the slower (1-Hz) update rate of the guidance system when compared to the inner-loop flight control system (50 Hz).

The performance of the dynamic inversion control law can be more fairly evaluated by comparing the body-rate commands to

the actual body rates. Figure 8 makes this comparison. The actual body rates ideally should look like those of the model or low-pass filtered command signals according to Eq. (17). It can be seen that the actual roll and yaw rates follow the commands very closely, which indicates that the inversion is nearly perfect in these two axes. Small short-term differences appear when comparing commanded pitch response to the actual response, which indicates that some pitch-axis modeling information is inaccurate. The differences are slight, however, and do not significantly impact the attitude tracking performance, as can be seen in Fig. 7. The appearance of high-frequency noise in the pitch response is caused by the modeling error in the control effectiveness matrix and null-space injection. Integration of an online system identification algorithm will reduce the modeling error and allow the control allocation to adapt to control effector failures or damage.




Fig. 9 Euler angle tracking with power pack out at 72 s.

Figure 9 shows the commanded and actual Euler angles for an ascent trajectory, where one of the turbopumps fails causing a 50% loss of engine thrust. This result was obtained without modification to the baseline control law other than reducing the differential throttle commands by 50%. Note that the system performs quite well with no large transients at the time of the failure.

The baseline control law operates quite well over the entire ascent trajectory, which covers a very wide range of flight conditions. The vehicle is launched vertically and accelerates to approximately Mach 9.5 at 180,000 ft. The vehicle center of gravity moves approximately 4m as the mass of the vehicle changes from 285,000 to 79,000 lbm in about 3 min. Note that this control law does not use gain scheduling and does not require linear models of the vehicle over the flight envelope for synthesis. The control law simply requires access to the nonlinear vehicle model parameters and estimates of the vehicle states as the flight progresses. The control design would be better described as model scheduled because a vehicle model is required and the control law is explicitly written in terms of the model parameters. This approach has a distinct advantage over traditional gain-scheduling techniques because a series of tuned classical or modern control designs does not have to be performed if the vehicle model is improved or a new trajectory must be flown. New vehicle configurations can be flown without changing the control system architecture or generating a new gain schedule.

Only the new a priori estimates of the model parameters must be loaded into the flight control law.

VI. Online System Identification

In the event of control effector damage or failures, a static approach to online system identification proposed by Chandler et al.⁵ and Buffington et al.⁹ is used to estimate the vehicle's changing control derivatives. The newly updated control derivatives are then used by the dynamic inversion control law to track the attitude guidance commands. The static identification approach provides a direct, non-iterative solution. Prior information about the system, such as the a priori knowledge of the effectiveness of the control surfaces, can be included in the identification algorithm in the form of stochastic constraints for better estimates of these values. The control derivative estimates can also be improved by lengthening the data window.

For the example in this paper, a control surface failure occurs when the X-33's entire right flap δ_1 is lost at time equal to 20 s, rendering its control derivatives in the pitch, roll, and yaw axes zero. The static identification method is used to estimate the control derivatives of the remaining effectors in all three axes. Without loss of generality, the following discussion will focus on the estimation of the roll control derivatives $\dot{p}_{\delta_1}, \dots, \dot{p}_{\delta_{12}}$ of the X-33's 12 control effectors $\delta_1, \dots, \delta_{12}$. This procedure given in the roll axis can then be extended to consider the remaining pitch and yaw axes.

The modeled roll acceleration equation of motion in the stability axis is

$$\dot{p} = \dot{p}_\beta \beta + \dot{p}_r r + \dot{p}_p p + \dot{p}_{pq} pq + \dot{p}_{\delta_1} \delta_1 + \dots + \dot{p}_{\delta_{12}} \delta_{12} + w_p + \text{higher-order terms} \quad (27)$$

The rolling acceleration coefficients due to side slip β , yaw rate r , roll rate p , and pitch rate q are \dot{p}_β , \dot{p}_r , \dot{p}_p , and \dot{p}_q , respectively. The roll acceleration coefficients due to the 12 control effectors $\delta_1, \dots, \delta_{12}$ are $\dot{p}_{\delta_1}, \dots, \dot{p}_{\delta_{12}}$, respectively. The measurement noise is w_p .

To estimate the new control derivatives, we first remove contributions of the side slip, pitch rate, roll rate, yaw rate, and the higher-order terms from Eq. (27):

$$\begin{aligned} \dot{\hat{p}} &= \dot{p} - [\dot{p}_\beta \beta + \dot{p}_r r + \dot{p}_p p + \dot{p}_{pq} pq + \text{higher-order terms}] \\ &= \dot{p}_{\delta_1} \delta_1 + \dots + \dot{p}_{\delta_{12}} \delta_{12} \end{aligned} \quad (28)$$

Concatenating k sampled measurements in Eq. (28), we have a data window of length k . Setting $n = 12$, the number of control effectors, we have

$$\underbrace{\begin{bmatrix} \dot{\hat{p}}_1 \\ \vdots \\ \dot{\hat{p}}_k \end{bmatrix}}_Z = \underbrace{\begin{bmatrix} \delta_{11} & \delta_{21} & \dots & \delta_{n1} \\ \vdots & \vdots & \ddots & \vdots \\ \delta_{1k} & \delta_{2k} & \dots & \delta_{nk} \end{bmatrix}}_H \underbrace{\begin{bmatrix} \dot{p}_{\delta_1} \\ \vdots \\ \dot{p}_{\delta_n} \end{bmatrix}}_\Theta + \underbrace{\begin{bmatrix} w_{p1} \\ \vdots \\ w_{pk} \end{bmatrix}}_W \quad (29)$$

Compactly, Eq. (29) can be written as

$$Z = H\Theta + W$$

where Z is a $k \times 1$ vector of measured roll accelerations due to the control effectors. H is a $k \times n$ regressor matrix of measured control surface deflections. The $n \times 1$ vector Θ is the rolling moment coefficient to be estimated. W is the system sensor noise. The $n \times 1$ vector W is the stochastic process of zero mean with the covariance $R(\theta) = r(\theta)I_k$. The minimum-variance estimate⁵ $\hat{\Theta}_{mv}$ of Θ is then

$$\hat{\Theta}_{mv} = (H^T R^{-1} H)^{-1} H^T R^{-1} Z$$

The standard of deviation of the estimate $\hat{\Theta}$ is then

$$\hat{\sigma} = \sqrt{(\tilde{Z}^T \tilde{Z}) / (k - n)}$$

where $\tilde{Z} = Z - H\hat{\Theta}$ is the return difference. The corresponding covariance P_{mv} of the estimate $\hat{\Theta}_{mv}$ is

$$P_{mv} = \hat{\sigma}^2 (H^T H)^{-1}$$

A priori knowledge of parameters to be estimated, such as the values of rolling moments of the unfailed aircraft obtained from experimental tests, can be used at this point to obtain a mixed estimate of the minimum-variance estimate and a priori values. The a priori values $\Theta_{apriori}$ with their associated covariance Q become the stochastic constraint on the final mixed estimate $\hat{\Theta}_{me}$ and its covariance:

$$\hat{\Theta}_{me} = \hat{\Theta}_{mv} + P_{mv}(P_{mv} + Q)^{-1}(\Theta_{apriori} - \hat{\Theta}_{mv})$$

$$P_{me} = [I - P_{mv}(P_{mv} + Q)^{-1}]P_{mv}$$

In calculating the mixed estimate $\hat{\Theta}_{me}$, the $k \times n$ moving data window H is updated by replacing the earliest values of the control surface deflections with their latest values. The same is done for the $k \times 1$ vector Z of the roll accelerations. Finally, the mixed estimate results $\hat{\Theta}_{me}$ are low-pass filtered (first order with 15-rad/s bandwidth) to smooth out the final results.

In the following example, we apply the preceding identification method to estimate the control derivatives of the X-33's 12 control effectors with a complete loss of the right flap δ_1 at 20 s. The vehicle has 12 control effectors to generate the pitching, rolling, and yawing moments. These control effectors consist of eight aerodynamic control surfaces: left and right body flaps δ_1 and δ_2 , left and right rudders

δ_3 and δ_4 , left and right inboard elevons δ_5 and δ_6 , as well as left and right outboard elevons δ_7 and δ_8 . The vehicle's propulsion-based control effectors include the left and right top engine quadrants δ_9 and δ_{10} and the left and right bottom engine quadrants δ_{11} and δ_{12} , respectively. Because of their large size and locations, the body flaps are the most effective of the eight aerocontrol surfaces. The engine quadrants, however, are the most effective of all 12 control effectors. The uncertainties associated with the control effectiveness parameters are directly reflected in their covariances. The uncertainties can depend on the flight conditions and the quality of the experimental data. The covariance of the effectors' a priori rolling accelerations are chosen to be $r(\delta_1) = r(\delta_2) = 0.01$ and $r(\delta_3) = \dots = r(\delta_8) = 0.001$ and $r(\delta_9) = \dots = r(\delta_{12}) = 0.0001$. Because of their large effectiveness relative to the other aerocontrol surfaces and the airflow interaction on their large surfaces, the covariance of the a priori estimates of the body-flap control effectiveness parameters are chosen to be larger than those of the less effective aerocontrol surfaces. The covariances associated with the engine quadrants are chosen to be smaller than those of the aerocontrol surfaces because the propulsive forces generated by the engine are less uncertain.

Figure 7 shows the nominal performance of the vehicle under no failure. With the loss of the right flap and without the online system identification, the controller fails to adapt to the changing rolling moments; thus, it is unable to track the guidance commands, as seen in Fig. 10. Under failure and with online system identification, one can see that the controller is able to track the guidance commands. The value of the mixed estimate of the right flap rolling control derivative is shown in Fig. 11. It is observed that rapid variation of the right flap's control derivative estimate is caused by the change in the slope of the roll command. The sensitivity of the right flap's estimate to the roll command can be reduced by choosing a slower filter. Furthermore, at lower dynamic pressure, the aerocontrol effectors $\delta_1, \dots, \delta_8$ are less effective. This occurs for time greater than 150 s. The rolling accelerations produced by the aeroeffectors are, thus, small in Eq. (28), whereas the engine differential throttle effectors $\delta_9, \dots, \delta_{12}$ are large. This difference in the effectiveness among the vehicle's effectors causes wide fluctuations in the estimates of the aeroeffectors' control derivative.

In Fig. 10, the right flap failure occurs 20 s into flight, and the vehicle begins to depart at approximately 110 s. The ability of the vehicle to remain stable for the next 70 s with a missing right body flap and no system identification can be attributed to the following two reasons.

First, the propulsion-based effectors play dominant roles in maintaining the vehicle's ability to track the guidance command, although the aerocontrols are also important as evidenced in the departure of the failed vehicle later in the flight. Initially, any increases in the tracking errors due to the missing body flap are feedback and compensated by the other control effectors, mainly the four engine quadrants. When only the nominal information about the failed vehicle is used, the control allocation scheme is able to distribute the control deflection commands. Stability and good tracking performance are then achieved as long as the remaining control effectors are not saturated.

Second, as fuel is expended during the ascent flight, the vehicle's center of gravity migrates toward the front of the vehicle causing a change in the vehicle dynamics. Consequently, the aerocontrol effectors become more effective in controlling the vehicle due to their lengthening moment arms. The stabilizing roles of the aerocontrol effectors, therefore, become more important. Without the knowledge of the failed vehicle dynamics and its control effectiveness, the control allocation algorithm fails to distribute properly the commanded control forces and moments among the remaining functional control effectors, resulting the departure of the vehicle.

In Fig. 11, the parameters being estimated are the rolling moments $L_{\delta_{flap}}$ generated by the left and right body flaps. They are fully dimensionalized and not normalized by the dynamic pressure \bar{q} . The a priori values of the roll derivatives are the outputs of a polynomial fit of the experimental data that depend on the vehicle's angle of attack, sideslip angle, altitude, and velocity. The roll control derivative of the left body flap is underestimated compared to its a priori values. The roll control derivative of the missing right




Fig. 10 Attitude tracking performance with lost right flap.




Fig. 11 Estimates of left and (failed) right flaps roll control derivatives (identification on at 20 s).

body flap fluctuates about the zero value as expected. Exact matching of the roll control derivatives to the a priori value for the left body flap and zero for the missing right body flap may be difficult to obtain because these parameters vary within the data window. These estimation errors are compensated by the robustness of the dynamic inversion control law. The overall control law is shown in Fig. 12.

There are three classes of actuation failures being considered:

1) In the locked actuator class, the failed actuator remains in a fixed position regardless of the command input to the device. As a result, constant pitching, rolling, and yawing moments are induced. Furthermore, the regressor matrix is singular because the matrix column corresponding to the locked actuator is zero.

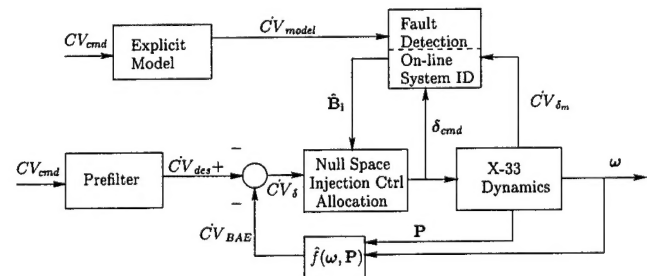


Fig. 12 Overall control law.




Fig. 13 Fault detection logic.

2) In the floating actuator case, the actuator position floats with the angle of attack of the vehicle. Its control effectiveness is, thus, zero. A loss of power or hydraulic pressure is the main cause of such a failure.

3) The third class is damaged arosurface. Although the actuator is in proper working order, damage to the vehicle aerocontrol surface is sustained. A fractional loss of the effector's surface area results in a proportional reduction in its control effectiveness.

The algorithm used to detect the three actuation failure scenarios contains the following parameters:

1) Response error thresholds are used for actuation failure detection. The response of the vehicle such as its pitch, roll, and yaw accelerations are constantly monitored and compared with those of the ideal, onboard model of the vehicle. Under an actuation failure, the magnitudes of response differences increase and exceed the preset error thresholds. Small error thresholds are used when the vehicle model is of high fidelity, whereas large error thresholds are necessary to avoid false alarms when the vehicle model is inaccurate.

2) A second parameter on which the fault detection algorithm depends, response error persistency, is the uninterrupted time duration during which the response errors exceed their thresholds. The fidelity of the ideal model determines how we set the persistency threshold for the response error. Longer persistencies of the response errors are used when the vehicle model contains greater uncertainties. In this way, false alarms can be further avoided.

3) The third parameter is the time length of the data window. Once the response errors have exceeded their thresholds and persisted for the required time, null-space injection is activated. Vehicle response and its actuator position data are then collected for use in the identification algorithm. The length of the moving data window determines how many data points will be collected. System identification takes place only after the moving data window is full. A temporal separation between the time when null-space injection is activated and the time when the system identification estimates the vehicle parameters is enforced. The temporal separation is equal to the time length of the data window.

When the three conditions are not met, such as during nominal flight conditions or after system estimates have been made and the damaged vehicle response errors have been brought below the threshold, the most recent mixed estimate of control effectiveness matrix is used. Figure 13 details the detection algorithm architecture.

The primary purpose of the detection algorithm is to avoid unnecessary computations and minimize control usage by activating null-space injection and system identification only when a fault occurs.

VII. Conclusions

A control architecture for the X-33 has been presented to demonstrate the feasibility of applying adaptive/reconfigurable control

technologies developed for fighter aircraft to RLVs. A baseline dynamic inversion control law with null-space injection control allocation demonstrated good attitude tracking capability over the entire ascent trajectory. This baseline control architecture is well suited for the insertion of direct and indirect adaptive control elements to improve the fault tolerance of these high-speed flight vehicles. Dynamic inversion allows designers to modify the control law easily when vehicle configurations change or when more accurate aerodynamic data become available. The designer simply updates the model parameters (tables, curve fits, etc.), and the control law synthesis is complete. Stability and robustness analyses should still be performed; however, a completed redesign of a traditional trajectory-dependent gain-scheduled launch vehicle control law is not necessary. The resulting control law can reduce mission design time because it is trajectory independent. Although the control law was demonstrated for the ascent flight phase only, the technique is general enough to cover the full flight envelope of such vehicles. The online system identification module has been integrated into the simulation and has been tested on a limited basis.

References

- Ward, D., and Barron, R., "A Self-Designing Receding Horizon Optimal Flight Controller," *Proceedings of the American Control Conference*, IEEE Publications, Piscataway, NJ, 1995, pp. 3490-3495.
- Brinker, J. S., and Wise, K. A., "Nonlinear Simulation Analysis of a Tailless Advanced Fighter Aircraft with Reconfigurable Flight Control Law," *Journal of Guidance, Control, and Dynamics*, Vol. 24, No. 5, 2001, pp. 903-909.
- Eberhardt, R. L., "Indirect Adaptive Flight Control of a Tailless Fighter Aircraft," *Proceedings of the 1999 AIAA Guidance, Navigation, and Control Conference*, AIAA, Reston, VA, 1999, pp. 466-476.
- Buffington, J. M., "Modular Control Law Design for the Innovative Control Effectors (ICE) Tailless Fighter Aircraft Configuration 101-3," TR AFRL-VA-WP-TR-1999, Air Force Research Lab., Wright-Patterson AFB, OH, 1999, pp. 93, 94.
- Chandler, P. R., Pachter, M., and Mears, M., "System Identification for Adaptive and Reconfigurable Control," *Journal of Guidance, Control, and Dynamics*, Vol. 18, No. 3, 1995, pp. 516-524.
- Schumacher, C., "Adaptive Flight Control Using Dynamic Inversion and Neural Networks," AIAA Paper 99-4068, Aug. 1999.
- Hall, C. E., Gallaher, M. W., and Hendrix, N. D., "X-33 Attitude Control System Design for Ascent, Transition, and Entry Flight Regimes," AIAA Paper 98-4411, Aug. 1998.
- Battin, R. H., *An Introduction to the Mathematics and Methods of Astrodynamics*, AIAA, New York, 1987, pp. 93, 94.
- Buffington, J. M., Chandler, P., and Pachter, M., "Integration of Online System Identification and Optimization-Based Control Allocation," AIAA Paper 98-4487, 1998.
- Petersen, J. A., and Bodson, M., "Fast Control Allocation Using Spherical Coordinates," *Proceedings of the 1999 Guidance, Navigation, and Control Conference*, AIAA, Reston, VA, 1999, pp. 1321-1330.